



GAU 277# 21# 4/6/01 JF
A33405 PCT USA - 072854.0112

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Howarth et al.
Serial No. : 09/601,695
Filed : August 7, 2000
For : SCHEDULING MEANS FOR DATA
SWITCHING APPARATUS

Group Art Unit 2774

RECEIVED
OCT 18 2000
TC 200 MAIL ROOM

CLAIM FOR PRIORITY UNDER 35 U.S.C. § 119

Assistant Commissioner for Patents

Washington, D.C. 20231

Sir:

A claim for priority is hereby made under the provisions of 35 U.S.C. § 119 for the above-identified PCT application based upon Great Britain application 9803301.2 filed February 18, 1998, and International Application PCT/GB99/00405 filed February 9, 1999.

Respectfully submitted,

Ronald B. Hildreth
Patent Office Reg. No. 19,498

(212) 408-2544
Attorney for Applicant

Baker Botts L.L.P.
30 Rockefeller Plaza
New York NY 10112

THIS PAGE BLANK (USPTO)



The
Patent
Office

PCT/GB 99 / 0 0 4 0 5

REC'D 03 MAY 1999	
WIPO	PCT

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP9 1RH

GB99/00405

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

PRIORITY DOCUMENT


SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

Signed

Dated

G D Court.
13th April 1999

THIS PAGE BLANK (US-10)



The
Patent
Office

18FEB98 E339014-1 003020
P01/7700 25.00 - 9803301.2

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

M98/0106/GB

2. Patent application number

(The Patent Office will fill in this part)

9803301.2

18 FEB 1998

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Power X Limited
Stafford Court
145 Washway Road
Sale
Manchester M33 7PE

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

6803233002

4. Title of the invention

Scheduling Means for Data
Switching Apparatus

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

McNeight & Lawrence
Regent House, Heaton Lane
Stockport, Cheshire SK4 1BS

Patents ADP number (if you know it)

0001115001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form	
Description	6
Claim(s)	1
Abstract	1
Drawing(s)	1

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature	Date
J G Lawrence	17.2.98

12. Name and daytime telephone number of person to contact in the United Kingdom
- J G Lawrence
0161 480 6394

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

SCHEDULING MEANS FOR DATA SWITCHING APPARATUS

This invention relates to scheduling means for data switching apparatus for use in computer-controlled digital data switching systems.

Many types of data transmission apparatus are known, all having their own particular features and systems. In all cases the intention is to allow data switching and transmission to be achieved as rapidly as the apparatus will allow. It is common for data to be sent in "packets" consisting of a predetermined number of bits of data plus control information indicating certain parameters of the data or its mode of transmission.

In data switching apparatus having an number of users there may be a requirement at any one time to set up a number of different interconnections between input ports and output ports. In any form of switch there is a limit to the number of simultaneous interconnections that may be formed. The switch is operating at its greatest efficiency when the greatest possible number of interconnections is formed and switching apparatus frequently includes what may be termed "scheduling means" in order to achieve this maximum number of interconnections.

A good scheduling scheme needs to balance the potentially conflicting objectives of making sure that all output ports are connected where there are requests for a connection to that port (efficiency), that high priority traffic is serviced quickly (prioritisation) and that low priority traffic is not ignored (fairness). The present invention addresses all of these issues and may, for example, be used with the data switching apparatus described and claimed in our co-pending British patent application No. 9717412.2

Various types of scheduling means are known. For example, United States Patent No. 5,500,858 describes one form of scheduling means in which requests for interconnections are considered and satisfied using what are called "rotating priority iterative matching desynchronising scheduler units". The "priority" in this case refers to priority given to input and output ports at any given time so as to ensure that each port has a fair chance of having a connection requests satisfied. The U.S. Patent goes on to describe how the scheme could be extended to handle requests at multiple priority levels but the scheme described would lack fairness, that is low priority requests would be ignored

under heavy load conditions where only higher priority requests would be satisfied.

5 It is an object of the present invention to provide data switching apparatus which includes scheduling means operable to satisfy a greater number of requests for interconnections than has previously been possible under such circumstances.

10 According to the present invention there is provided scheduling means for data switching apparatus having a plurality of input ports and a plurality of output ports between which data having one of a predetermined number of priority levels is to be passed, which scheduling means includes a first pipeline stage operable to satisfy at least some of the requests for interconnections which are applied to the scheduling means, a priority mixer to which are applied those requests for interconnections which were not satisfied by the first pipeline stage together with requests of different priority levels and operable to select which of those requests should be further considered, and at least one further pipeline stage to which are applied said further requests and operable to satisfy such of those requests as are possible and were not satisfied by any preceding pipeline stage.

20 The present invention overcomes the problems associated with the known prior art by using existing types of scheduling units, (for example those described in U.S. Patents Nos. 5,500,858 and 5,267,235, though any scheduling means which operates as described herein may be used) but connecting them in a novel arrangement. The scheduling means to be described may, for example, be used with the data switching apparatus described and claimed in our co-pending British patent application No. 9717412.2.

30 The invention will now be described with reference to the accompanying drawing, which shows a schematic block diagram of one embodiment of the invention.

35 The drawing shows three pipeline stages 10 to 12, with a priority mixer 13 connected between pipeline stages 10 and 11. Input connections and output connections are provided to the various pipeline stages and the priority mixer as shown and the operation of the arrangement will be described below.

Each of the pipeline stages 10, 11 and 12 operates to receive input connection Request Vectors RV_i at Priority level P_i and a

Connection Vector CVi. In response to these inputs the pipeline stage generates output signals Queue Return QRet, Request Vector out RVo, Priority out Po and Connection Vector out CVo. The Request Vectors are bit fields where each bit corresponds to a possible connection between one of the input ports and one of the output ports of the data switching apparatus. That is, if there are n input ports and m output ports, the Request Vectors will be $n \times m$ bits wide, where a bit that is set indicates that a connection is being requested from the corresponding input port to the corresponding output port, whilst a bit that is clear indicates that such a connection is not being requested at this time. The Priority fields Pi and Po indicate the priority of the connection being requested at input (RVi) and output (RVo) respectively. The connection requests from each input port are all of the same priority, though the connections requested from different input ports may be of different priorities. The Connection Vector signal CVo defines connections which are to be made by a switching matrix (not shown). They indicate which input port, if any, is to be connected to each output port of the data switching apparatus. The Queue Return signals Qret represent connection requests that cannot be satisfied. These requests are returned to the input queues of the data switching apparatus ready to be requested again. The operation performed by each pipeline stage is to consider the connection requests at RVi and satisfy as many of them as possible, adding details of each satisfied connection to any already present at CVi and presenting the combined set of connections at CVo. Since each input port and each output port may only be involved in one connection at any given time, any connection requests which involve an input port or output port which is already part of a satisfied connection request may no longer be satisfied within the present set of connections and so such requests are returned to the input queues (signal Qret), to be considered as part of a subsequent set of connections. The remaining connection requests (those for which the corresponding input and output ports are still available for connection) are presented at the RVo output of the pipeline stage in order for them to be considered by a subsequent pipeline stage. Any such requests at the output of the last pipeline stage 13 (where there is no subsequent pipeline stage to consider them) are returned to the input queues of the data switching apparatus, as with the Qret output.

Consider now the overall operation of the scheduling means described above. At the input to the first pipeline stage 10 requests for connections RV_i at a single priority level P_i are presented. The first pipeline stage 10 then attempts to satisfy as many of these requests as possible. By choosing P_i to be a high priority more frequently than choosing it to be a lower priority, high priority packets are delivered on average more quickly than lower priority packets. If there is only a small number of requests at the priority P_i then the first pipeline stage 10 will not make many connections and most of the input and output ports will not be utilised within the set of connections created by this stage, nor will there be many connection requests outstanding at that priority level which may be satisfied by the remaining pipeline stages 11 and 12. For this reason, the priority mixer 13 is introduced between the first and second pipeline stages 10 and 11. Applied to priority mixer 13 are connection requests RV_{2i} of priorities other than P_i , the priorities of the requests being denoted by P_{2i} . The priority mixer 13 decides, for each input port, whether to pass on to the second pipeline stage 11 the requests RV_i remaining at priority level P_i from the first pipeline stage 10 or the new requests RV_{2i} . The decision is made on the basis of choosing whichever set of requests has the highest number of requests that could still be satisfied within the current Connection Vector CV_i , taking into account which input and output ports are already used by satisfied connection requests. This leads to higher connectivity within the data switching apparatus than if only requests of a single priority were considered, that is it is more efficient. It also allows good performance for low priority traffic in the absence of any higher priority traffic, since the low priority requests may be presented at the second pipeline stage 11 via the priority mixer 13, regardless of how infrequently low priority requests are selected to be presented to the first pipeline stage 10. The Connection Vector output CV_o of the last pipeline stage 13 defines all the connections that are to be made by the data switching apparatus in the next cycle of operation.

Since mixing requests of different priorities at the second and subsequent stages of the scheduling means leads to greater efficiency, it might seem like a good idea to do the same at the first pipeline stage 10. However, the scheduling units that make up each pipeline stage

do not themselves take any account of the priority of each request, that is they treat all requests equally. Hence if requests of different priorities were presented to the first pipeline stage there would be no concept of priority at all within the scheduling means, since the low
5 priority requests would compete equally with the high priority requests for connections. Thus it will be seen that the first pipeline stage provides prioritisation and fairness as defined above, whilst the subsequent pipeline stages provide efficiency.

Further pipeline stages may be added if it is felt that three
10 stages are not able to provide sufficiently high efficiency of switch utilisation. There is a trade-off between switch utilisation (how many connection requests may be satisfied at any time) and latency (each pipeline stage takes time to operate), so that the number of pipeline stages required will depend upon what balance of these factors is
15 required for a particular embodiment. In general, more pipeline stages are needed to create maximal sets of connections as the number of ports in the data switching apparatus increases. In addition, priority mixer elements may be placed between others of the pipeline units to further increase efficiency if desired.

Instances may occur when, for example, certain connections between
20 input ports and output ports are to be retained for a period of time (static or permanent connections). Alternatively, it may be necessary at certain times to block specified input ports or output ports (during a period of system maintenance, for example). Similarly, it may be
25 necessary at certain times to set up connections from one input port to more than one output port at the same time (multicast). All of these facilities may be incorporated into data switching apparatus which uses the scheduling means described above. This is done by connecting appropriate logic to the inputs of the first pipeline stage 10 and/or
30 the priority mixer 13. For instance, to create a permanent connection between an input port and an output port, the CVi input of the first pipeline stage 10 could be preset to indicate the required connection(s) rather than having all of its bits clear (indicating no pre-existing connections). Input and output ports may be blocked by masking off the
35 appropriate bits of the RVi input of the first pipeline stage 10 and the RV2i input of the priority mixer 13. Multicast connections may be made in the same way as permanent connections, except that more than one

- 6 -

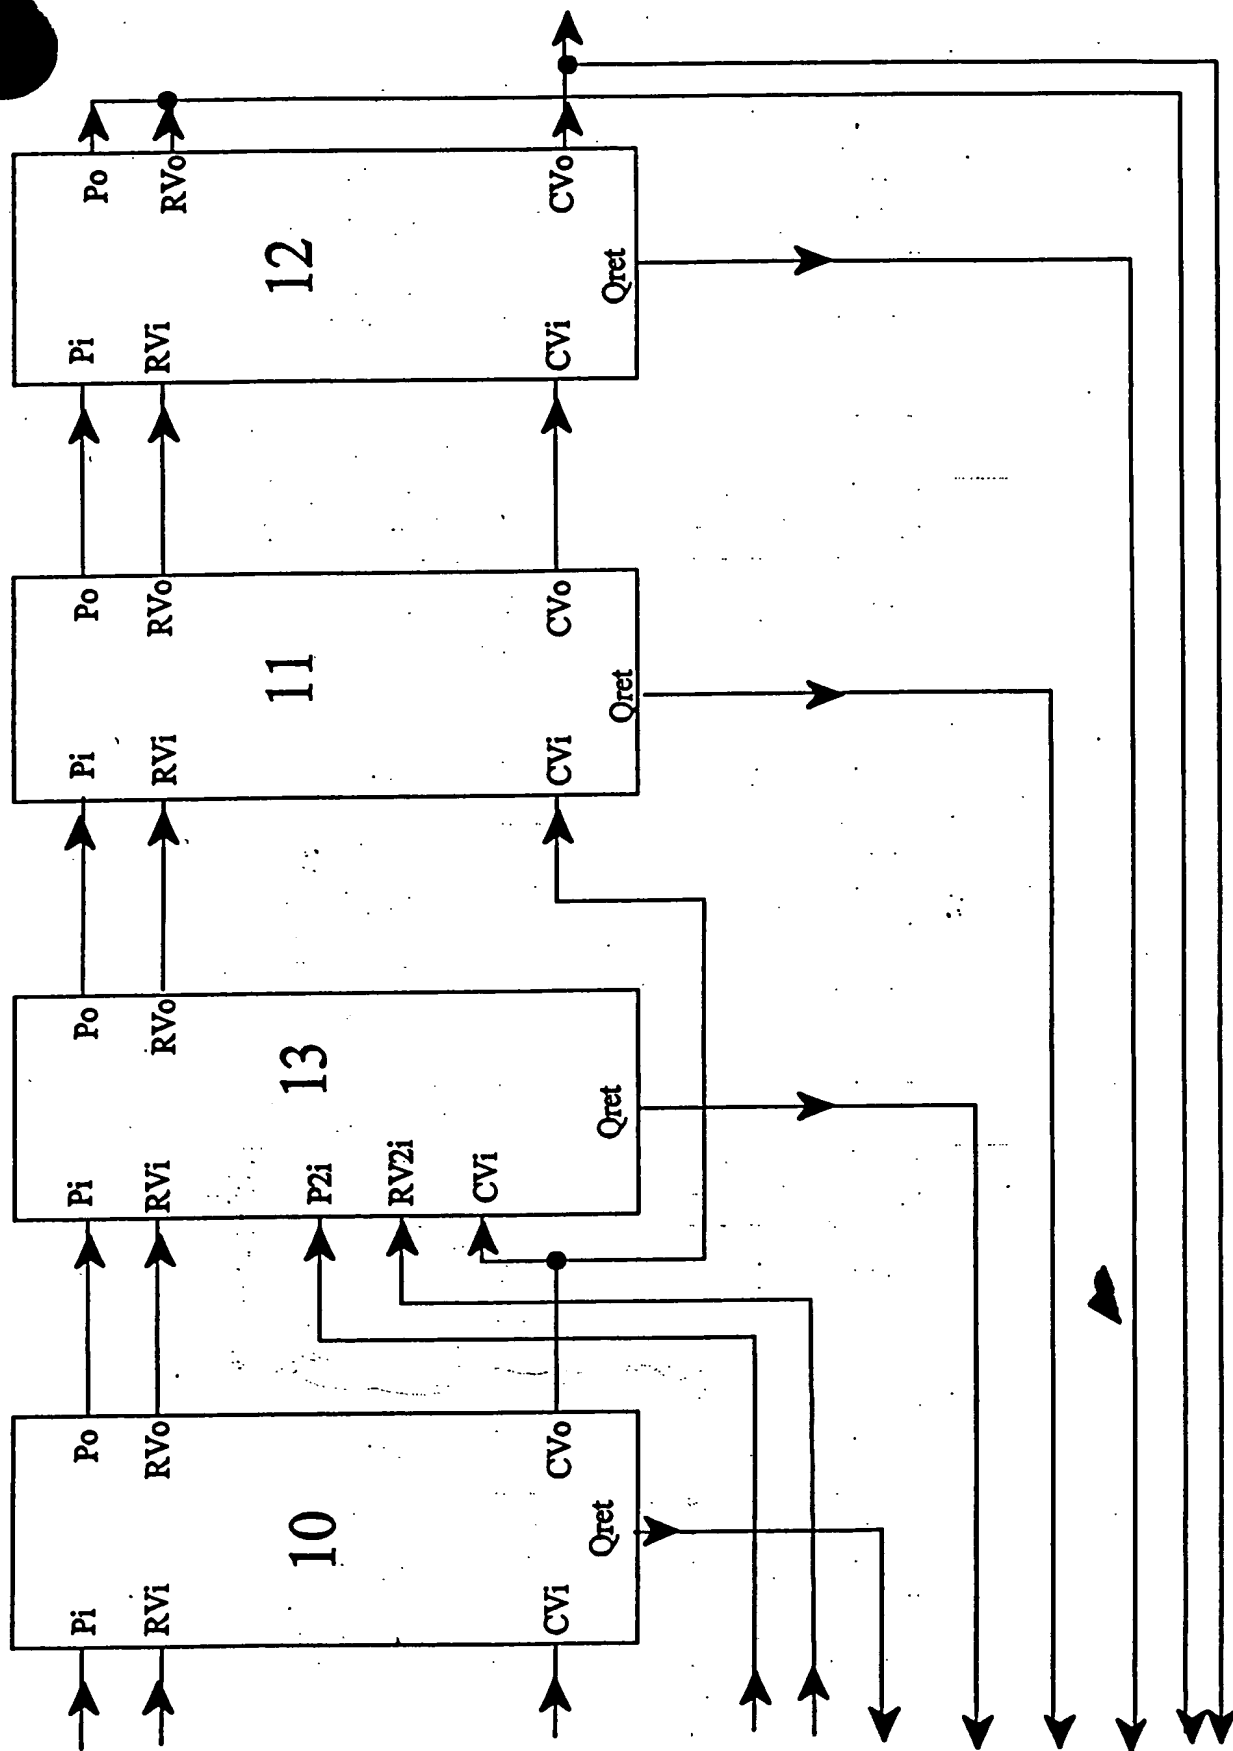
output port is set up to be connected to the desired input port. In all of these cases, the scheduling means works around the existing connections or blocked ports, making whatever connections it can between the remaining unconnected and non-blocked input and output ports.

CLAIMS

1. Scheduling means for data switching apparatus having a plurality of input ports and a plurality of output ports between which data having one of a predetermined number of priority levels is to be passed, which scheduling means includes a first pipeline stage operable to satisfy at
5 least some of the requests for interconnections which are applied to the scheduling means, a priority mixer to which are applied those requests for interconnections which were not satisfied by the first pipeline stage together with requests of different priority levels and operable to select which of those requests should be further considered and at
10 least one further pipeline stage to which are applied said further requests and operable to satisfy such of those requests as are possible and were not satisfied by any preceding pipeline stage.
2. Scheduling means as claimed in Claim 1 in which requests for interconnections applied to the first pipeline stage are all of the same
15 priority level.
3. Scheduling means as claimed in either of Claims 1 or 2 in which priority mixers are inserted between any pairs of the pipeline stages.
4. Scheduling means as claimed in any one of Claims 1 to 3 in which
20 different types of predetermined connections may be set up between some of the ports at the input to the scheduling means such that the scheduling means is free to set up connections only between the ports not so connected.
5. Scheduling means as claimed in any one of the preceding claims in which requests for connections to or from any of the input or output
25 ports may be inhibited, thereby preventing any connections being made to or from those ports.
6. Scheduling means for data switching apparatus having a plurality of input ports and a plurality of output ports between which data having one of a predetermined number of priority levels is to be passed,
30 substantially as herein described with reference to the accompanying drawing.

ABSTRACT

Scheduling means is provided for data switching apparatus having a plurality of input ports and a plurality of output ports between which data having one of a predetermined number of priority levels is to be passed. The scheduling means includes a first pipeline stage (10) to which are applied requests for said interconnections from the input ports and operable to satisfy at least some of the requested interconnections. A priority mixer (13) is provided to which are applied those requests for interconnections which were not satisfied by the first pipeline stage together with requests of a different priority level. The priority mixer (13) operates to select which of those requests should be further considered and applied to a second pipeline stage (11). This is operable to satisfy such of those requests as are possible. Further pipeline stages (12) are provided, to which are applied those requests not satisfied by any preceding pipeline stages.



PCT/GB 99/00 405.

Beresford & Co.

16/3/99

THIS PAGE BLANK (USPTO)